A Rule-based Tool for Assisting Colormap Selection

Lawrence D. Bergman
bergman@watson.ibm.com

Bernice E. Rogowitz
rogowtz@watson.ibm.com

Lloyd A. Treinish
lloydt@watson.ibm.com

IBM Thomas J. Watson Research Center
Yorktown Heights, NY

Abstract

The paper presents an interactive approach for guiding the user's select of colormaps in visualization. PRAVDAColor, implemented as a module in the IBM Visualization Data Explorer, provides the user a selection of appropriate colormaps given the data type and spatial frequency, the user's task, and properties of the human perceptual system.

1: Introduction

Visualization is a process of mapping data onto visual dimensions to create a pictorial representation. A successful visualization provides a representation which allows the user to gain insight into the structure of the data, or to communicate aspects of this structure effectively [2] [4] [22]. Even with modern visualization systems, which give the user considerable interactive control over the mapping process, it can be difficult to produce an effective visualization. One strategy for improving this situation is to guide the user in the design of a visualization. In our previous work, we have described an interactive rule-based architecture for incorporating such guidance, and have described certain perceptual and cognitive rules which may be relevant [14] [15] [16].

In this paper, we focus on improving the user's selection of colormaps. To do so, we have built a library of colormaps, and a set of perceptual rules for selecting appropriate maps based on the structure of the data and the goal of the visualization. We have encapsulated this rule-based colormap selection process as a tool, *PRAVDAColor*, in the IBM Visualization Data Explorer software package, and demonstrate how this module can be incorporated into visualization applications involving the mapping of color onto two- and three-dimensional surfaces. This implementation demonstrates the viability of the technique, and provides a testbed for evaluating the rules.

I.1: Interactive Rule-Based Architecture

We have previously presented a rule-based architecture called PRAVDA (Perceptual Rule-Based Architecture for Visualizing Data Accurately) for assisting a user in making choices of visualization parameters [15] [16]. In this architecture, each visualization operation can be associated with rules which constrain the set of choices the user is offered. The architecture also provides for linkages between rules that control different visualization operations, with a choice of parameters for one operation constraining choices that are available for others. This network of linked, intelligent operations helps guide the user through the complex process of designing a visualization.

In our previous work, we have described the general principles for implementing such an assemblage of rule-based visualization operations. In this paper, we describe a full implementation of one of these operations, colormap selection. In *PRAVDAColor*, perceptual rules constrain the set of colormaps offered to the user based on system-provided metadata (data type, data range), metadata computed by algorithm (spatial frequency) and metadata provided by the user (the visualization task). This is in contrast to previous rule-based systems for visualization which do not explicitly support user tasks, color perception, or interactivity in the guidance they offer (e.g., [18]).

I.2: Rule-Based Colormap Selection - Limitations of Current Technology

Perhaps the most common operation in visualization is mapping the values of a variable onto a color scale. Despite the importance of this operation, the creation and selection of colormaps is often not adequately supported in modern visualization systems, which typically offer the user a default color map and a tool for creating custom colormaps. More importantly, however, these systems do not guide the user in selecting which

(See color plates, page CP-14)

1070-2385/95 $4.00 © 1995 IEEE

colormaps will help the user understand the structure of the data, segment the data meaningfully, or highlight important characteristics of the data.

The most common default colormap, the "rainbow" colormap, is a hue-based scale from blue, through a rainbow of colors, to red. When this scale is mapped onto scalar data, the user is conceptually mapping a linear scale in hue onto a scalar variable. Perceptually, however, this scale does not appear linear. Equal steps in the scale do not correspond to equal steps in color, but look instead like fuzzy bands of color varying in hue, brightness and saturation. When mapped onto scalar data, this colormap readily gives the user the erroneous impression that the data are organized into discrete regions, each represented by one of the rainbow colors. This can lead the user to infer structure which is not present in the data and to miss details that lie completely within a single color region [13] [15].

Some visualization systems provide the user with tools for creating alternatives to the default colormap. These custom colormap facilities allow the user to construct almost any conceivable colormap, but tend to be difficult to use for creating complex colormaps. This inadequacy has led us to develop our own colormap generation tool, which is described below. Even when an adequate tool is provided, however, deciding what colormap to create is essentially an unbounded problem. Some systems, such as Interactive Data Language, address this problem by providing a library of colormap lookup tables [9]. They do not, however, provide guidance about their application or alert the user to constructs that may introduce visual artifacts. In this paper, we improve on this approach by offering the user colormaps guided by principles of human perception. In this way, *PRAVDAColor* can offer the user a richer set of choices than a single default colormap, while ensuring that each choice is appropriate to the user's visualization task and the data being represented.

II: A Taxonomy for Colormap Selection

Table 1 shows our working taxonomy for the generation of colormaps. Following other researchers [5] [3] [24] [11] [12] [7], we find the data types described in measurement theory to be very powerful in characterizing data representation strategies. The data type is shown in the left-most column. Currently, we support interval and ratio data, but plan to extend this approach to include ordinal and nominal data types. For each data type, we distinguish between low and high spatial frequency data, depicted in the second column.

The next three columns contain recommendations on creating colormaps for these four data types which are appropriate for three different representation tasks. Following our earlier work [13] [15], an isomorphic task is one where the goal of the representation is to faithfully reflect the structure in the data. In a segmentation task, the goal is to divide the data into perceptually distinct categories. In a highlighting task, the goal is to call attention to particular features in the data.

The notes within each cell show our understanding to-date of the characteristics of colormaps which best meet these requirements. The following section describes the psychophysical rationale underlying these selections.

II.1: Isomorphic Representation

In order to accurately represent continuous data, the visual dimension chosen must appear continuous to the user. In an MRI image, for example, the degree of magnetic resonance, a continuous variable, is represented as a gray-level because continuous variations in gray-scale appear continuous to the user. That is, if the resonance increases monotonically over a spatial region, the brightness of the image will appear to increase monotonically over that spatial region. An inappropriate colormap, however, could create a visual representation which does not look monotonically increasing, for example, the default colormap described above.

Candidate colormaps which preserve the monotonic relationship between data values and perceived magnitude can be drawn from experiments done by Stevens [20]. Stevens identified a set of sensory dimensions (visual, auditory, and tactile) for which a monotonic increase in stimulus intensity produced a monotonic increase in perceived magnitude. In particular, he found the shape of this relationship to be a power law, with each sensory dimension characterized by the exponent of this power law. Perceived brightness obeys a power relationship with physical intensity (gray-scale) over a very large range of gray scales, making it a very good candidate for representing ratio or interval data. Another good candidate includes the color attribute saturation, although the dynamic range for saturation is significantly smaller than for brightness.

II.2: The Importance of Spatial-Frequency for Ratio and Interval Data

Mapping continuous variables onto perceptually continuous dimensions will only give a faithful representation of the structure of the data if the spatial characteristics of the representation are taken into account. Visual sensitivity to spatial variation differs for the hue and luminance mechanisms in human vision. The luminance mechanism is tuned to higher spatial

Data Type	Spatial Frequency	Representation Task: Isomorphic	Segmentation	Highlighting
Ratio (true zero)	Low	*Luminance:* uniform *Hue:* opponent or complementary pairs *Saturation:* monotonically increasing from gray	- Even number of segments - Many segments OK	Larger range for highlighted features
	High	*Luminance*: monotonically increasing *Hue:* opponent or complementary pairs *Saturation:* monotonically increasing from gray	- Even number of segments - Fewer segments	Smaller range for highlighted features
Interval	Low	*Luminance:* uniform *Hue*: opponent pairs *Saturation:* monotonically increasing from gray	- Many segments OK	Larger range for highlighted features
	High	*Luminance:* monotonically increasing *Hue*: uniform or small hue variation *Saturation:* monotonically decreasing	- Fewer segments	Smaller range for highlighted features

Table 1. A Taxonomy of Colormaps based on Data Type, Representation Task, and Principles of Perception

frequencies (that is, high resolution, finely detailed, or small-grained features). Colormaps which include a luminance component, therefore, can adequately represent high spatial frequency information. The hue mechanism is tuned to lower spatial frequencies. Thus, saturation-based colormaps, which display variations in the magnitude of a hue, would be inadequate for conveying high spatial frequency information, but well-suited for representing larger-scale spatial variations [17].

To ensure appropriate choices, the colormap rules in *PRAVDAColor* use metadata about spatial frequency to constrain the set of selectable colormaps. Metadata about spatial frequency are not generally included in data sets, so in *PRAVDAColor* it is computed. If the metadata indicate that the data, when realized on the display screen, would produce high spatial-frequency variations, then the rules offer the user colormaps which contain a monotonically-varying luminance component. These would include simple gray-scale colormaps and also colormaps which also vary in hue (e.g., maps from dark green to light green, or maps from dark cyan to light lavender).

If the data contain predominantly low spatial frequencies, information about their spatial variation is more effectively carried by the hue channel, and the rules would offer the user colormaps with variations in saturation within a single hue or variations in saturation across two hues, such as the opponent pairs [6] blue/yellow (saturated blue, through isoluminant gray, to saturated yellow) and green/red (saturated green, through isoluminant gray, to saturated red).

Data containing both fine spatial detail and gradual variations across space could be mapped to colormaps which combine these characteristics, such as dark saturated green to light saturated red.

PRAVDAColor incorporates these psychophysical relationships into its rule set. To create a faithful (isomorphic) representation of ratio and interval data, *PRAVDAColor* recommends colormaps which vary in luminance and saturation for high and low spatial frequency data respectively. A special distinction is made for ratio data where it is important for the representation to preserve the existence of a zero point. *PRAVDAColor* treats this as two separate scales, one above and one below the zero, and uses hue to distinguish them.

II.3: Colormaps for Segmentation Tasks

Some of the principles for providing isomorphic colormaps for ratio and interval data are also effective in creating maps for segmented data. The luminance component conveys monotonicity in the data for high spatial frequency data, while the saturation component can be used to convey monotonicity in low spatial-frequency data. Since the steps are explicitly defined, however, luminance steps can also be effectively used for low spatial-frequency data. In creating a segmented colormap, it is important that the segments are each discriminably different from one another, which limits the number of steps which can be represented. We have found that more steps can be effectively discriminated for low spatial-frequency data than for high. For ratio data, where the zero is a semantically-important attribute, *PRAVDAColor* only offers segmented maps with an even number of steps (with a transition at the zero level).

II.4: Colormaps for Highlighting Tasks

Rules for selecting colormaps which highlight particular features in the data can be drawn from the literature on attention (e.g., [21]). Based on these rules, *PRAVDAColor* offers colormaps which allow the user to identify ranges of data to highlight perceptually. We currently provide colormaps which identify the mid-range data value, or set a threshold and highlight data values which exceed it.

III: Interface and Implementation

III.1: The *PRAVDAColor* Module

PRAVDAColor has been implemented as a module for the IBM Visualization Data Explorer (Data Explorer) software package [1]. Data Explorer facilitated the development of *PRAVDAColor*'s capabilities by providing a unified, extensible data model and polymorphic modules.

Figure 1 shows a sample Data Explorer data-flow network and the visualization that it produces. The *PRAVDAColor* module accepts as input the data set to be visualized (read by the *Import* module in this example) and a user representation task index (created by the interactive *Selector* module).

Figure 2 shows the same data represented with the default hue-based colormap (from the *AutoColor* module). This colormap is less successful in preserving the spatial structure of the data.

The *PRAVDAColor* user is presented with a set of colormaps appropriate to the data set and specified representation task. These colormaps are presented in a panel which pops up on the display when *PRAVDAColor* is executed as shown in Figure 1. Each colormap is displayed as a colorbar and has an associated button, used for selecting that colormap. *PRAVDAColor* automatically selects one of the colormaps as an initial output, which is used to color the data. The user is then free to examine a sequence of colormaps applied to the data by simply clicking the virtual button under each color bar.

In addition to allowing interactive selection, *PRAVDAColor* allows the user some control over how the colormap is to be mapped onto the data set. Range sliders allow the minimum, maximum and midpoint values of the data to be assigned to particular color values.

The *PRAVDAColor* tool consists of two main portions. The first is a macro that selects the set of colormaps to be presented to the user based on characteristics of the data and a user-specified representation task. This macro, called *ColorMapLookup*, examines two characteristics of the data -- spatial frequency, and presence or absence of a zero-crossing.

A simple spatial frequency analysis, consisting of computing the normalized standard deviation of the difference between the original (structured) data grid and a low-pass filtered grid, is used to classify data as either high or low spatial frequency. Presence or absence of a zero crossing is used to distinguish ratio from interval data. Based on these calculations and a user-specified representation task, *ColorMapLookup* reads an appropriate set of colormaps from a library into memory.

The second major component of *PRAVDAColor* is the interactive colormap displayer/selector. This component is implemented as an outboard module in Data Explorer called *ColorMapPicker*, written in C using Motif/Xlib. Outboard modules are separately compiled and linked executables which are invoked by Data Explorer and communicate with the Data Explorer executive (main computational component) via socket. As described above, *ColorMapPicker* allows the user to iteratively select colormaps, as well as dynamically alter the assignment of colormap values to data values.

III.2: Colormap construction

To construct colormaps, we developed a tool similar to the interactive tool described by Rheingans [10]. Figure 3 shows the 3-dimensional colormap tool, an image that has been created using its output, and a view of its output in the Data Explorer *Colormap* editor.

The colormap tool displays a hue-lightness-saturation (HLS) double-cone populated by color swatches. We chose to represent each swatch by a single square polygon rather than a solid to avoid the color artifacts created by a shading model.

The user can rotate and zoom the color space representation using the mouse. To ensure interactive updates, we draw only an outline of the space when the mouse is moving, filling in the color swatches when the mouse stops. By selecting color swatches with the mouse, a color path is constructed interactively.

The tool provides an 'undo' and 'move' function for editing the path. At any time, the user can send the current colormap to the rest of the Data Explorer network by clicking on a virtual button. The output colormap linearly interpolates between the selected swatches in HLS space. Similarly, the user may specify that a discrete colormap is to be produced containing equally-sized segments for each selected swatch with no interpolation. This is useful in creating colormaps for segmented representation tasks.

The colormap tool was implemented in C using Motif/Xlib for interface elements and event handling, and OpenGL to create the three-dimensional graphics. The module converts all colors from HLS, used internally, to hue-saturation-value (HSV), required by Data Explorer.

IV: Results -- Applications of *PRAVDAColor*

PRAVDAColor has been informally tested within IBM as a prelude to making these tools available for evaluation by interested users of Data Explorer. *PRAVDAColor* has been integrated into various operational visualization activities.

One application of *PRAVDAColor* is to preserve the spatial structure of data. Figures 1 and 2 show the fluid density from a simulation of the noise produced by a jet aircraft engine. The representation in Figure 2 is dominated by the segmentation inherent in the default rainbow type of colormap. Figure 1 shows *PRAVDAColor* being used to select an isomorphic colormap. It shows more of the fine spatial structure and turbulence inherent in these continuous data.

Figure 4 illustrates how different tasks result in different colormaps suggested by *PRAVDAColor*. This figure shows two views of the result of a photochemical grid model of transport and deposition of airborne pollutants over the midwestern portion of the United States on June 26, 1987 at 18:00 local time. Ozone pollution concentration is shown in parts per billion by volume (ppbv). The isomorphic colormap (left) supports the task of capturing the inherent dynamics of the model by showing a snapshot of atmospheric motion (e.g., roughly circular filaments in yellow corresponding to higher ozone concentrations). The segmented colormap (right) supports the task of isolating regions. Higher pollution levels (e.g., above 160 ppbv) are clearly visible as yellow and red over Lake Michigan, to the east of Chicago. It should also be noted that this colormap allows the user to see some artifacts of the limited grid resolution of the model.

The application of *PRAVDAColor* is not confined to two-dimensional data. Figure 5 shows one time step of a regional, three-dimensional weather model shown as an isosurface of wind speed at 20 m/sec. The surface is colored by the temperature values in the computed volume interpolated on the isosurface. *PRAVDAColor* is used to create an isomorphic colormap corresponding to the full temperature volume.

PRAVDAColor can easily be incorporated into actual applications built with Data Explorer. Figure 6 is a screen dump of a generalized application that provides cartographic representations from a selection of available parameters stored in a user-defined data set. The package was extended by replacing the standard *Colormap* editor with *PRAVDAColor*. This enhanced application allows a user to select a colormapping task, for which *PRAVDAColor* then offers a set of choices. The image shows a segmented colormap applied to filled contours of total column ozone displayed as orthographic maps for the northern and southern hemispheres. The seasonal ozone depletion is visible as a black region over the south pole for these data taken on October 1, 1991. The available choices are illustrated in the *ColorMapPicker* panel.

PRAVDAColor can be used whenever a default colormap would be used, and can be used simultaneously for several operations, including, for example, two- and three-dimensional objects, contours, cutting planes, and streamlines. Since *PRAVDAColor* offers the user several choices for colormapping each, it can be quite easy to avoid some of the artifacts commonly introduced by color mixing. A future direction for this work is to provide feedback between operations as a way of further assisting the user.

V: Discussion

For complex colormap selection activities, *PRAVDAColor* has reduced the length of time required to develop satisfactory results. This has been sufficiently encouraging to warrant further enhancements, and an effort to make the tools available to current Data Explorer users interested in evaluating them and providing feedback for their improvement.

V.1: Importance of Spatial Frequency in Colormap Selection

Previous guidance for colormap selection has established the importance of various classes of colormaps appropriate for specific visualization tasks (e.g., [7] [24]). However, there can be considerable variation in the collection of colormaps that effectively satisfy a given task from one data set to the next. The analysis of the spatial frequency coupled with a large characterizing the appropriate representation of the data. The user is shielded from the intricacies of psychophysical research and color theory, and is automatically offered colormaps with significant luminance variations for representing fine spatial structure. Similarly, the user is offered colormaps with significant variation in saturation for representing coarse spatial structure.

Our simple algorithm for computing spatial frequency based on the differences between the low-pass and original version of the data is obviously limited. For instance, we have simply chosen to classify the data as "low" or "high" spatial frequency. A finer-grain scale may prove useful, as well as a method for characterizing data which has both coarse and finer spatial structure.

Having demonstrated the importance of spatial frequency, it may be useful to pursue the computation in the frequency domain. This would entail filtering the Fourier-transformed data to characterize various distributions of spatial frequency information, in two and three dimensions.

The frequency-domain-based methodology, however, has inherent limitations. Since it assumes regularly gridded data, it would not be useful for curvilinear, partially regular, multizone or irregular data. It is also less efficient than the spatial domain approach we have adopted, and is not sensitive to disparate scaling in each dimension of a grid. Selecting an appropriate, efficient, method for calculating spatial frequency, and evaluating the use of these metadata to characterize the data for the rules is a direction for ongoing research.

V.2: Integration of Intelligence into Existing Systems

In our earlier work, we established the importance of interactive rule-based tools to extend the utility of modern visualization software systems. These rules allow the user to capitalize on a broad range of visualization capabilities while minimizing the number of iterations required to create more appropriate and better visualizations. As a first step in this effort, the implementation of *PRAVDAColor* has demonstrated that this is feasible and practical. We expect other visualization systems to be extensible to incorporate *PRAVDAColor*-like tools through the addition of custom code as well as existing visual programming modules (e.g., [8] [19] 23]).

V.3: Extensions

In this paper, we have presented a Data Explorer module for guiding the user in selecting colormaps. The next step would be to add additional rule-based operations and provide feedback between the operations.

The PRAVDA architecture can be readily extended to allow the user to interactively modify the metadata and the rules. For example, the user could choose to modify the algorithm for computing spatial frequency, or decide to see what mapping choices were available if the goal of the visualization were changed. Likewise, the rules could be changed to reflect new insights or information. For example, if the user realized that the structure being sought was confined to a certain range of values in the data set, the rule could be modified so that data in this range would be highlighted by mapping it to an appropriate visual dimension. Our first focus has been on implementing rules based on perceptual principles, but the approach could be easily extended to include other types of rules, for example, rules about mappings or conventions particular to certain domains [17].

VI: Conclusions

A framework for rule-based guidance can improve the effectiveness of systems by assisting the user in making two types of appropriate representation choices. One concerns domain-independent factors, such as ensuring that data content is faithfully represented in images and that perceptual artifacts are not erroneously interpreted as data features. A second type concerns task-dependent factors. For example, different advice on representation is required depending on whether the goal of visualization is exploration or presentation.

The PRAVDA architecture explicitly incorporates guidance based on principles of human perception, cognition and color theory [16]. These principles are incorporated in rules which the user can select during the visualization process. Depending on the higher-level characteristics of the data, the rule constrains the way in which the data are mapped onto visual dimensions. This architecture has been utilized to build a tool which eases the burden of creating colormaps for many visualization applications.

The *PRAVDAColor* module has been implemented as an initial proof-of-concept of the PRAVDA rule-based architecture for advising a visualization developer. We have developed a set of rules that simplify the task of colormap selection. The user is presented a set of choices that are appropriate, based on characteristics of the data being visualized and the desired form of the representation.

Using this rule-based colormap advisor, we have constructed a set of visualizations that demonstrate the advantages of colormaps that conform to perceptual principles over an uninformed choice. We anticipate users will find the current implementation helpful in constructing visualizations, with the utility increasing as we add additional rules and intelligent visualization operations to the system.

VII: Acknowledgments

We thank the Data Explorer development team for their cooperation in this effort and valuable discussion that led to the successful implementation of *PRAVDAColor*.

Data shown in Figures 1 and 2 are available courtesy of Combustion Research and Flow Technology. Data shown in Figures 3 and 6 are available courtesy of NASA/Goddard Space Flight Center. Data shown in Figure 5 are available courtesy of Glenn Wightwick, IBM Australia. Data shown in Figure 4 are available courtesy of the US Environmental Protection Agency, Research Triangle Park, NC.

VIII: References

[1] Abram, G. and L. Treinish. "An Extended Data-Flow Architecture for Data Analysis and Visualization". To be published in **Proceedings of the IEEE Computer Society Visualization '95**, October 1995.

[2] Bertin, J. **Semiology of Graphics**, 1967, reprinted by University of Wisconsin Press, Madison, WI, 1983.

[3] Carswell, C. M. and C. D. Wickens. "The Perceptual Interaction of Graphical Attributes: Configurality, Stimulus Homogeneity, and Object Integration". **Perception & Psychophysics**, pp. 157-168, 1990.

[4] Cleveland, W.S. **The Elements of Graphing Data,** Monterey, California: Wadsworth, Inc., 1991.

[5] Della Ventura, A. and R. Schettini, "Computer Aided Color Coding". **Communicating with Virtual Worlds**, N.M. Thalmann and D. Thalmann (eds.), Springer-Verlag, Tokyo, pp. 62-75, 1991

[6] Hurvich, L. M. and Jameson, D. "Some quantitative aspects of an Opponent-Colors Theory: II. Brightness, Saturation, and Hue in Normal and Dichromatic Vision." **J. Optical Soc. of America**, 45, pp. 602-616, 1955.

[7] Lefkowitz, H. and G. T. Herman. "Color Scales for Image Data". **IEEE Computer Graphics and Applications**, 12, n. 1, pp. 72-80, January 1992.

[8] Rasure, J. and C. Wallace. "An Integrated Data Flow Visual Language and Software Development Environment". **Journal of Visual Languages and Computing**, 2, pp. 217-246, 1991.

[9] Research Systems, Inc. **IDL Basics**, Interactive Data Language Version 3.5, November, 1993.

[10] Rheingans, P. and B. Tebbs. "A Tool for Dynamic Exploration of Color Mappings". **Computer Graphics** 24/2, pp. 145-146, 1990.

[11] Robertson, P. K. "Visualizing Color Gamuts: A User Interface for the Effective Use of Perceptual Color Spaces in Data Displays". **IEEE Graphics and Applications**, pp. 50-63, September 1988.

[12] Robertson, P. K, M. Hutchins, D.R. Stevenson, S. Barass, C. Gunn and D. Smith, "Mapping Data into Colour Gamuts: Using Interaction to Increase Usability and Reduce Complexity". **Computer Graphics** (UK), v. 18, no. 5, pp. 653-665, September/October 1994.

[13] Rogowitz, B. E., D. T. Ling and W. A. Kellogg. "Task Dependence, Veridicality, and Pre-Attentive Vision: Taking Advantage of Perceptually-Rich Computer Environments." **Proceedings of the SPIE**, Volume 1666, **Human Vision, Visual Processing and Digital Display**, pp. 504-513, 1992.

[14] Rogowitz, B. E. and D. A. Rabenhorst. "CRAFT: A Tool for Customizing Color and Font Selections Guided by Perceptual Rules". **Proceedings of the IS&T and SID Color Imaging Conference**, November, 1993.

[15] Rogowitz, B. E. and Treinish, L. A. Data Structures and Perceptual Structures. **Proceedings of the SPIE/SPSE Symposium on Electronic Imaging**, Volume 1913, February 1993, 600-612.

[16] Rogowitz, B. E. and Treinish, L. A. An Architecture for Perceptual Rule-Based Visualization. **Proceedings of the IEEE Computer Society Visualization '93 Conference**, pp. 236-243, October 1993.

[17] Rogowitz, B. E. and Treinish, L. A. Using Perceptual Rules in Interactive Visualization. **Proceedings of the SPIE/SPSE Symposium on Electronic Imaging**, Volume 2179, February 1994, 287-293.

[18] Senay, H. and Ignatius, E., "A Knowledge-Based System for Visualization Design", **IEEE Computer Graphics and Applications**, 14, no 6, pp. 36-47, November, 1994.

[19] Silicon Graphics Computer Systems. "IRIS Explorer". **Technical Report BP-TR-1E-01** (Rev. 7/91).

[20] Stevens, S. S. "Matching Functions Between Loudness and Ten Other Continua," **Perception and Psychophysics**, 1, pp. 5-8, 1966.

[21] Treisman, A. and Gelade, G. "A Feature Integration Theory of Attention". **Cognitive Psychology**, 18, pp. 643-662, 1980.

[22] Tufte, E. **The Visual Display of Quantitative Information**, Connecticut: Graphics Press, 1983.

[23] Upson, C., T. Faulhaber, D. Kamins, D. Laidlaw, D. Schlegel, J. Vroom, R. Gurwitz and A. van Dam. "The Application Visualization System: A Computational Environment for Scientific Visualization". **IEEE Computer Graphics and Applications**, 9, n. 4, pp. 30-42, July 1989.

[24] Ware, C. "Color Sequences for Univariate Maps: Theory, Experiments and Principles." **IEEE Computer Graphics and Applications**, 8, n. 5, pp.41-49, 1988.

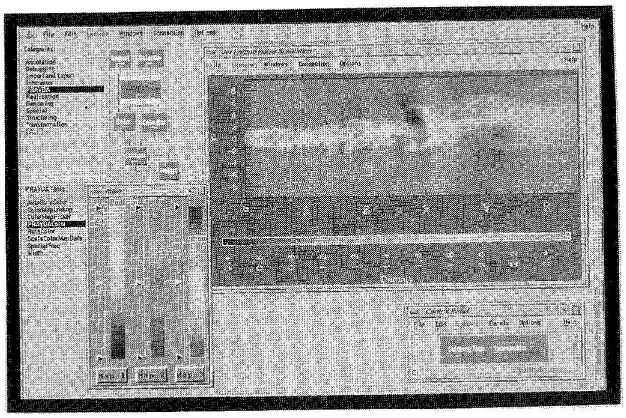

Figure 1. Data Explorer Visual Program Incorporating *PRAVDAColor*.

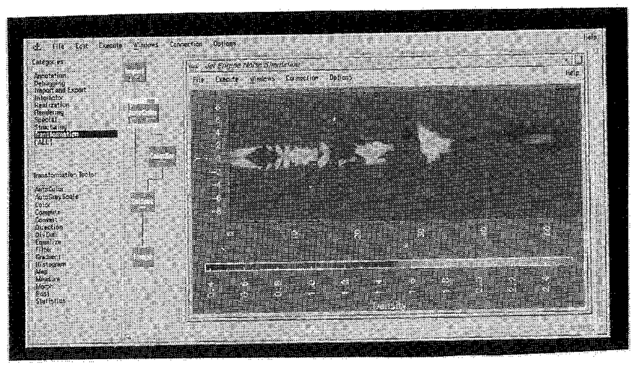

Figure 2. Data Explorer Visual Program Using the Default Colormap Tool (*AutoColor*).

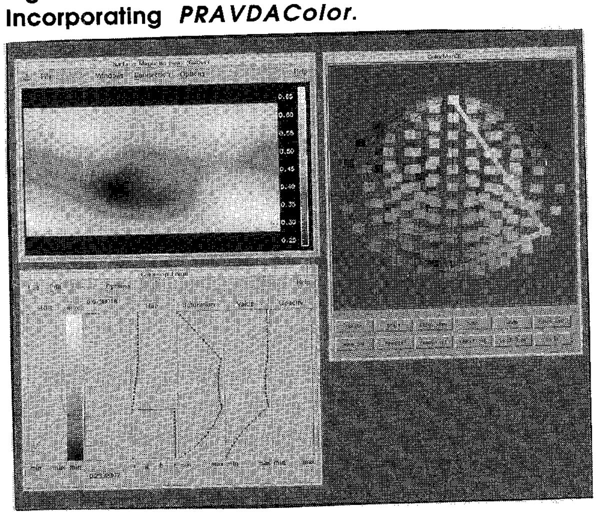

Figure 3. Construction of a Colormap with a Three-Dimensional Colormap Tool.

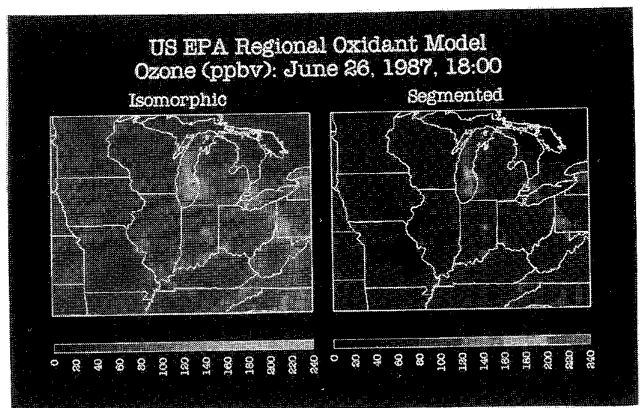

Figure 4. Photochemical Pollution Model with Isomorphic and Segmented Colormaps.

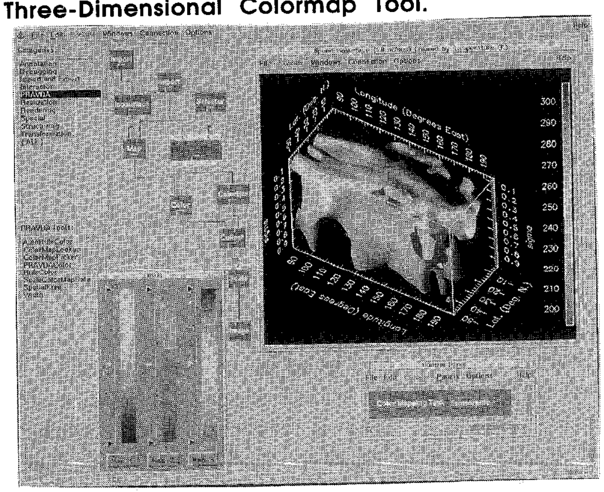

Figure 5. Data Explorer Visual Program Using *PRAVDAColor* with Three-Dimensional Data.

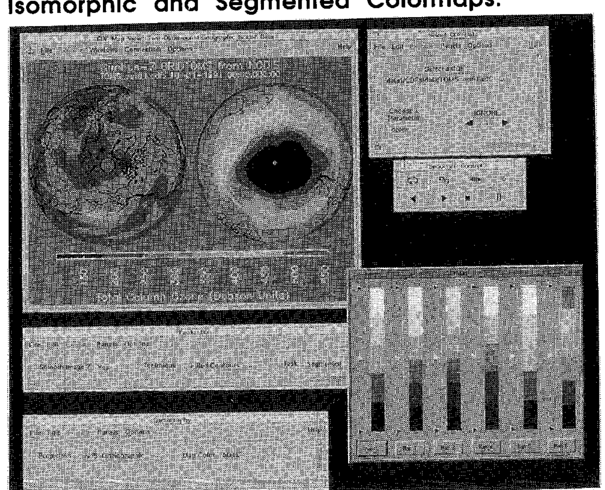

Figure 6. Data Explorer Application Incorporating *PRAVDAColor*.